

Jet Propulsion Laboratory  
California Institute of Technology

# An Overview of the Mars 2020 Perseverance Rover's Enhanced Path-Planner


*Workshop on Planetary Exploration Robots*

*IROS 2020, Online*

*October 29, 2020*

Olivier Toupet

Robotic Aerial Mobility Group Supervisor  
Mobility and Robotic Systems Section, JPL





**Mars 2020 Project**

# What I Do At JPL



Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav








# Why a New AutoNav for M2020?



- MSL AutoNav (GESTALT) is too conservative to handle the obstacle-rich terrain expected at Jezero Crater:
  - Treats the rover as a 5m diameter disk
    - 2.2m larger than true vehicle width
    - Makes it **impossible to traverse 15% CFA** terrains
  - **Low frequency terrain undulations** (that can be traversed) within each inscribed disc **are indistinguishable from obstacles** which leads to false positives
- M2020 mission relies heavily on AutoNav:
  - Nearly 75% of our drive distance between ROIs will be done with AutoNav
  - Each sol: 50m of blind driving + 144m of AutoNav on average



- Traverse rate
  - 100m/h in both benign and complex terrains
  - Average cycle time  $\leq 36$ s for 1m steps
- Translates to the following metrics:
  - Success rate
    - $\geq 90\%$  in benign terrains
    - $\geq 75\%$  in complex terrains
  - Path inefficiency
    - $\leq 15\%$  in benign terrains
    - $\leq 35\%$  in complex terrains
- Canonical landing site slope and CFA distribution:

Benign terrain  
Complex terrain

CFA Slope	0 - 7%	7 - 10%	10 - 12%	12 - 15%
15 - 20°	2%	1.5%	1%	0.5%
10 - 15°	5%	4%	1%	0.5%
5 - 10°	20%	10%	1%	0.5%
0 - 5°	40%	10%	2%	1%


# ENav Planner Components



Jet Propulsion Laboratory  
California Institute of Technology


M2020 AutoNav

## Global Planner



- Gives cost from the end of tree to goal
- Routes computed on 200m x 200m map
- 1 m resolution
- Considers slope, roughness, KOZ, KIZ


## Local Planner



- Selects best path for the next 6m from finite # of options

## ACE

(Approx. Clearance Est.)




- Evaluates clearance using approx. kinematics
- Run every 25cm
- Also checks tilt, wheel drop, etc

# Approx. Clearance Evaluation (ACE)



Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav




- At sampled poses along candidate arcs, ACE checks that:
  - Suspension and attitude angles are within bounds
    - $13^\circ / 30^\circ$  for rocker / bogie angles with  $10^\circ / 25^\circ$  reactive check
    - $30^\circ$  for roll and pitch
    - $30^\circ$  for tilt with  $25^\circ$  reactive check
  - Wheels don't drive over unknown terrain in near field
  - Belly pan clearance  $\geq 25\text{cm}$
  - Wheel drop height  $\leq 35\text{cm}$
- Evaluations are always conservative
  - Wheel drop: difference between max and min heights over the **entire** wheel box
  - Clearance computed as difference between lowest belly point and highest terrain point **anywhere** under rover belly
  - Assumes wheels may settle on the lowest terrain point

# ENav Planner Timeline



- Process Disparity
  - Updates the 2.5D heightmap based on latest stereo mesh & rover position
- Analyze Terrain
  - Updates the costmap based on the heightmap, KOZs, rover position, and distance to goal
- Select Arc
  - Selects the next arc to drive based on the heightmap, costmap, KOZs, rover pose, and goal
- Evaluate Arc
  - Evaluates safety of arc based on KOZs and terrain






- Thinking-While-Driving requires predicting the future pose of rover at the start of the plan
  - Pose uncertainty due to slip accumulated since last VO
- Some paths may require close proximity to surrounding obstacles (i.e. rover may straddle rocks)
  - Slip prediction accuracy impacts both path performance (feasibility and efficiency) and rover safety
- Unexpected slip resulting in deviations from planned path is the main safety concern
  - Must avoid unsafe terrain which could cause reactive safety faults, large wheel drops, or high-centering by enlarging the ACE wheel boxes

# Accounting for Slip



- Main idea: expand ACE footprints to account for max slip
  - Expansion is proportional to prediction distance (distance since last measurement)
    - Prediction distance is smaller for rotational than translational slip (more frequent IMU measurements than VO updates)



# ENav Slip Model Overview




Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav

- ENav slip model includes:
  - Translational slip
    - Omnidirectional
    - Downslope
  - Yaw slip
- Robust and conservative: expand ACE wheel boxes so safety conditions (belly pan clearance, wheel drop, etc) are met with predicted slip


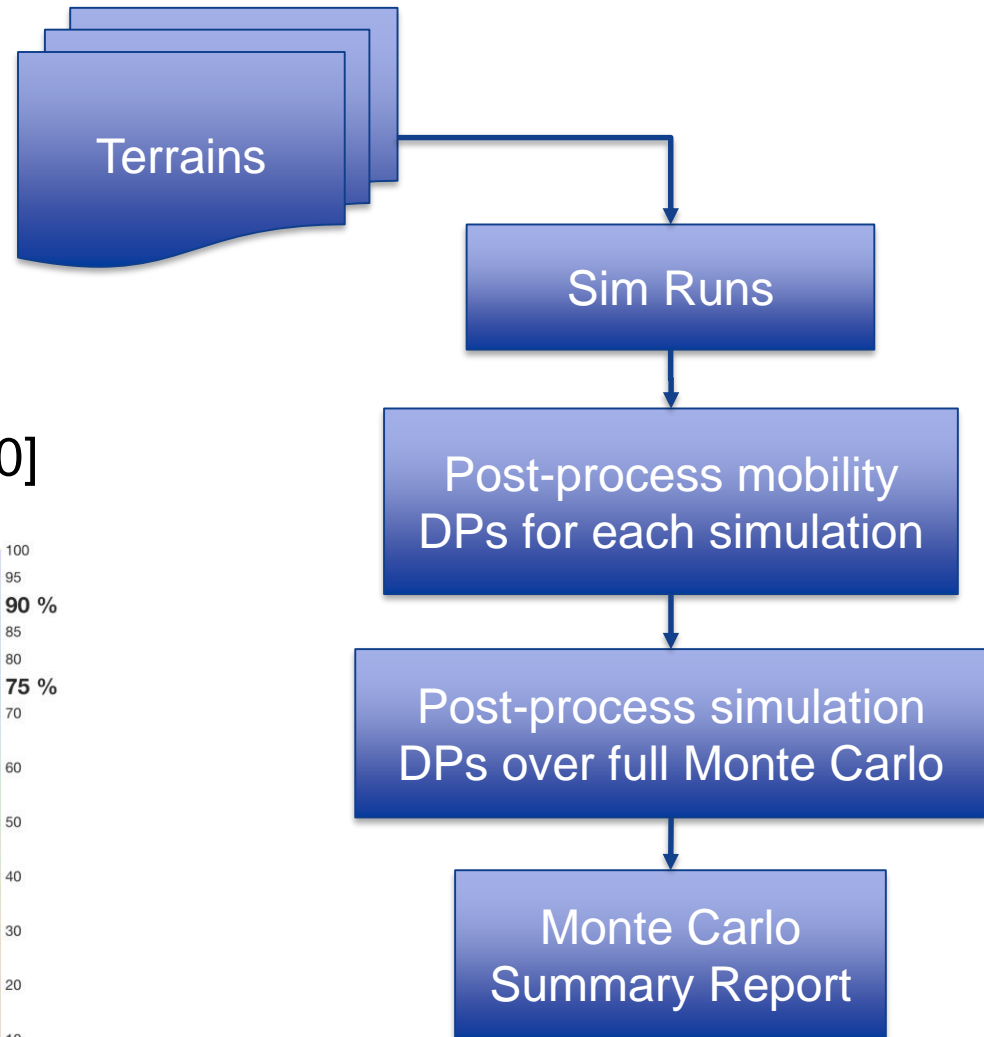
## 1.3 Omnidirectional slip



- Exhaustive Monte Carlo simulation testing with automated metrics assessment reports
- Comprehensive flight software unit tests
- Nightly integrated simulation runs (e.g. long multi-sol drives)
- Hardware-in-the-loop testing in the Mission System and Flight Software Testbeds (MSTB / FSWTB)
- Periodic (but less frequent) field tests in the Mars Yard:
  - With the Scarecrow surrogate testbed
  - Now with the Vehicle System Testbed (VSTB)



- Terrain Parameters
  - Slope Magnitudes (deg):
    - [0, 5, 10, 15, 20]
  - Slope Headings (deg):
    - [0, 45, 90, 135, 180]
  - CFAs (%): [0, 7, 10, 12, 15, 20]








# Monte Carlo Testing Results




Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav




# Scarecrow Testing Video




Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav





# Questions?



Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav







# Back Up

# Belly Pan Clearance Budget



Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav






# Wheel Box Tuning with MSL Data



Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav

- Tune ACE to be sufficiently conservative with respect to slip observed on MSL
- Justifications: similar vehicle design, real Mars data, statistically significant number of samples
- Assumption: AutoNav will not be used on sandy terrain
  - Justification: according to MTTT analysis, ~1% of distance on strategic route will be on sandy terrains in both JEZ and NES
  - Why we need this assumption: ENav agnostic to terrain type; tuning wheel box sizes for sandy terrain is overly conservative





# Results of MSL Slip Study




Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav

- Tuned ACE wheel boxes to contain the rover wheels for all 6,683 arcs
- Omnidirectional slip parameters:
  - Forward: 0
  - Backward: 10% (25cm minimum)
  - Inward: 0
  - Outward: 25%
- Yaw slip standard deviation: 2.7deg/m
- Extra lateral margin of 10cm
  - on each side of each wheel box
- Conservative downslope slip table




- We mitigate the risk of under-estimating slip by:
  - Using short planning steps
  - Using a conservative slip model when planning
    - Takes into account both translational and rotational (yaw) slip
    - informed by past data collected on Mars and most recent slip measurements during the drive
  - Frequent evaluations of currently executing arc
    - *MOM* will check path safety based on latest yaw measurement every 0.25m
  - Extra 10cm of lateral margin on the ACE wheel boxes
    - Provides robustness to sudden yaw slip of up to 4deg (which would have to occur within 0.25m)
    - Max yaw slip observed on MSL within 25cm = 2.5deg



## Local Height Map

- Dense height map around rover (5cm resolution, 15m radius grid)
- Produced from merging stereo vision meshes
- Used for calculating plane fits, belly clearance, wheel drop height, etc



## Global Cost Map



- Sparse cost map (1m resolution, 100m radius grid)
- For estimating cost to travel from end of local path to goal
- Cost includes roughness, slope, and penalty for unknown terrain






## Plane Fit Algorithm (2.5D Least Squares)

1. Anchor point  
(centroid of cloud to avoid numerical issues)
2. Estimate slope of plane in x and y direction  
(linear regression in two directions simultaneously by assuming each  $z_i = a \cdot x_i + b \cdot y_i + c$ )
3. Transform two slopes to 3D normal vector
4. Calculate statistics:
  1. roughness (Mean Squared Deviation from plane)
  2. Maximum deviation (furthest absolute deviation from plane)




# Local Planner: Tree




Jet Propulsion Laboratory  
California Institute of Technology

M2020 AutoNav


Depth 1: TIP




Depth 2: Arc



Depth 3: TIP




Depth 4: Arc



- Tree includes both arc *and* turn-in-place (TIP)
- Current tree configuration
  - Depth = 4
    - TIP-Arc-TIP-Arc
  - # of branches
    - 11-11-5-5
    - 3025 leaf nodes in total
  - Arc length: 3m
    - 6m total
  - Max turn angles
    - 172-150-90-150°
    - All parameters are configurable
- Special paths (explained later)
  - Retreat path
  - Heritage paths


- Add the previously selected path to the tree in the next iteration because it is likely to be feasible
  - Heritage paths (forward): Unexecuted portion + 1m extensions
  - Retreat path (backward): Executed portion

## Iteration N




- Add the previously selected path to the tree in the next iteration because it is likely to be feasible
  - Heritage paths (forward): Unexecuted portion + 1m extensions
  - Retreat path (backward): Executed portion

Iteration N+1



# Cost Function



Path cost =  
(entire path)

## Time to the end of tree

Includes time for:

- Driving
- Turning
- Steering

+

## Time from the end of tree to the goal

Comes from global planner

+

## Penalty

Includes:

- Backward path
- Tilt
- Roughness

ACE cost =  
(local path)

## Safety

• Inf if not meeting requirements on:

- Clearance
- Tilt
- Rocker/bogie angles
- Wheel drop

+


## Unknown Terrain

- Unknowns within 2m : Inf
- Unknowns beyond 2m : finite but heavily penalized

+

## Path Margin

- Penalizes obstacles within 1m laterally
- Nearer obstacles are penalized more heavily




# Local Planner: Path Selection



## Bottom line:

- Clearance evaluation (i.e., ACE) is expensive
- Rank paths based on *path cost*, and run ACE only on high-ranked paths






# Local Planner: Path Selection



## Bottom line:

- Clearance evaluation (i.e., ACE) is expensive
- Rank paths based on *path cost*, and run ACE only on high-ranked paths



## Algorithm:

1. Sort *all* the paths in a tree by path cost


ID	Path cost	ACE cost	Total cost
5	50.1		
2	50.3		
8	52.0		
6	53.0		
3	53.2		
1	55.6		
4	60.6		
9	62.1		
7	66.5		

# Local Planner: Path Selection



## Bottom line:

- Clearance evaluation (i.e., ACE) is expensive
- Rank paths based on *path cost*, and run ACE only on high-ranked paths




## Algorithm:

1. Sort *all* the paths in a tree by path cost
2. Run ACE on the top  $N$  paths

ID	Path cost	ACE cost	Total cost
5	50.1	Inf	Inf
2	50.3	5.1	55.4
8	52.0	Inf	Inf
6	53.0	0.0	53.0
3	53.2		
1	55.6		
4	60.6		
9	62.1		
7	66.5		


}  $N$

# Local Planner: Path Selection



## Bottom line:

- Clearance evaluation (i.e., ACE) is expensive
- Rank paths based on *path cost*, and run ACE only on high-ranked paths



## Algorithm:

1. Sort *all* the paths in a tree by path cost
2. Run ACE on the top  $N$  paths
3. If feasible paths are found, choose the “*best*” one among them

ID	Path cost	ACE cost	Total cost
5	50.1	Inf	Inf
2	50.3	5.1	55.4
8	52.0	Inf	Inf
6	53.0	0.0	53.0
3	53.2		
1	55.6		
4	60.6		
9	62.1		
7	66.5		

Selected


$N$

# Local Planner: Path Selection



## Bottom line:

- Clearance evaluation (i.e., ACE) is expensive
- Rank paths based on *path cost*, and run ACE only on high-ranked paths



## Algorithm:

1. Sort *all* the paths in a tree by path cost
2. Run ACE on the top  $N$  paths
3. If feasible paths are found, choose the “*best*” one among them
4. If no feasible path is found, keep going down the list and choose the first feasible path

ID	Path cost	ACE cost	Total cost
5	50.1	Inf	Inf
2	50.3	Inf	Inf
8	52.0	Inf	Inf
6	53.0	Inf	Inf
3	53.2	Inf	Inf
1	55.6	Inf	Inf
4	60.6	5.4	66.0
9	62.1		
7	66.5		

}  $N$


Selected

# Local Planner: Path Selection



## Bottom line:

- Clearance evaluation (i.e., ACE) is expensive
- Rank paths based on *path cost*, and run ACE only on high-ranked paths



## Algorithm:

1. Sort *all* the paths in a tree by path cost
2. Run ACE on the top  $N$  paths
3. If feasible paths are found, choose the “*best*” one among them
4. If no feasible path is found, keep going down the list and choose the first feasible path
5. If no feasible path is found at all, fail and replan

ID	Path cost	ACE cost	Total cost
5	50.1	Inf	Inf
2	50.3	Inf	Inf
8	52.0	Inf	Inf
6	53.0	Inf	Inf
3	53.2	Inf	Inf
1	55.6	Inf	Inf
4	60.6	Inf	Inf
9	62.1	Inf	Inf
7	66.5	Inf	Inf

}  $N$